

What Is Claimed Is:

1. A method to select features for maximum entropy modeling, the method comprising:
  - determining gains for candidate features during an initialization stage and for only top-ranked features during each feature selection stage;
  - ranking the candidate features in an ordered list based on the determined gains;
  - selecting a top-ranked feature in the ordered list with a highest gain; and
  - adjusting a model using the selected using the top-ranked feature.
2. The method of claim 1, wherein the gains of the candidate features determined in a previous feature selection stage are reused as upper bound gains of remaining candidate features in a current feature selection stage.
3. The method of claim 2, wherein the top-ranked feature is selected if its determined gain is greater than the upper bound gains of the remaining candidate features.
4. The method of claim 1, wherein the top-ranked feature is selected when a gain of the top-ranked feature determined using a currently adjusted model is greater than the gains of remaining candidate features determined using a previously adjusted model.
5. The method of claim 1, wherein gains for a predefined number of top-ranked features are determined at each feature selection stage.
6. The method of claim 1, further comprising:
  - re-evaluating gains of all remaining candidate features at a pre-defined feature selection stage.
7. The method of claim 1, wherein only the un-normalized conditional probabilities that satisfy a set of selected features are modified.

8. A method to select features for maximum entropy modeling, the method comprising:

- (a) computing gains of candidate features using a uniform distribution;
- (b) ordering the candidate features in an ordered list based on the computed gains;
- (c) selecting a top-ranked feature with a highest gain in the ordered list;
- (d) adjusting a model using the selected top-ranked feature;
- (e) removing the top-ranked feature from the ordered list so that a next-ranked feature in the ordered list becomes the top-ranked feature;
- (f) computing a gain of the top-ranked feature using the adjusted model;
- (g) comparing the gain of the top-ranked feature with a gain of the next-ranked feature in the ordered list;
- (h) if the gain of the top-ranked feature is less than the gain of the next-ranked feature, repositioning the top-ranked feature in the ordered list so that the next-ranked feature becomes the top-ranked feature and an order of the ordered list is maintained and repeating steps (f) through (g); and
- (i) repeating steps (c) through (h) until one of a quantity of selected features exceeds a predefined value and a gain of a last-selected feature falls below a predefined value.

9. The method of claim 8, wherein the step (f) of computing the gain of the top-ranked feature includes computing the gain of a predefined number of top-ranked features.

10. The method of claim 8, wherein the gains of all remaining features at a predefined feature selection are re-evaluated.

11. The method of claim 7, wherein gains of a majority of the candidate features remaining at each feature selection stage are reused based on a model adjusted in a previous feature selection stage.

12. A processing arrangement system to perform maximum entropy modeling in which one or more candidate features derived from a corpus of data are incorporated into a model that predicts linguistic behavior, the system comprising:

a gain computation arrangement to determine gains for the candidate features during an initialization stage and to determine gains for only top-ranked features during a feature selection stage;

a feature ranking arrangement to rank features based on the determined gain;

a feature selection arrangement to select a feature with a highest gain; and

a model adjustment arrangement to adjust the model using the selected feature.

13. The processing arrangement of claim 12, wherein feature ranking arrangement is configured to re-use gains of remaining candidate features determined in a previous feature selection stage using a previously adjusted model.

14. The processing arrangement of claim 12, wherein the gain computation arrangement and is configured to determine gains for top-ranked features in ascending order from a highest to lowest until a top-ranked feature is encountered whose corresponding gain based on a current model is greater than gains of the remaining candidate features.

15. The processing arrangement of claim 12, wherein the gain computation arrangement is configured to determine gains for a predefined number of top-ranked features at each feature selection stage.

16. The processing arrangement of claim 15, wherein the predefined number of top-ranked features is 500.

17. The processing arrangement of claim 11, wherein gains of all candidate features remaining at a predefined feature selection stage are re-evaluated.

18. A storage medium having a set of instructions executable by a processor to perform the following:

ordering candidate features based on gains computed on a uniform distribution to form an ordered list of candidate features;

selecting a top-ranked feature with a largest gain to form a model for a next stage;

removing the top-ranked feature from the ordered list of the candidate features;

computing a gain of the top-ranked feature based on a model formed in a previous stage;

comparing the gain of the top-ranked feature with gains of remaining candidate features in the ordered list;

including the top-ranked feature in the model if the gain of the top-ranked feature is greater than the gain of a next-ranked feature in the ordered list;

inserting the top-ranked feature in the ordered list so that the next-ranked feature becomes the top-ranked feature and an order of the ordered list is maintained, if the gain of the top-ranked feature is less than any of the gains of the next-ranked feature in the ordered list;

repeating the steps of computing the gain of the top-ranked feature, comparing the gains of the top-ranked and next-ranked features until the gain of the top-ranked feature exceeds the gains of ordered candidate features; and

terminating the method if one of a quantity of selected features reaches a pre-defined value and a gain of a last feature reaches a pre-defined value.